Ontology-Augmented Multi-Agent Reinforcement Learning for Enhancing EV Charging Network Recovery

Mohamed Abdelfattah[1]
[1]Technische Universität Berlin, Berlin, Germany
E-mail(s): m.abdelfattah.ibrahim.mohamed@tu-berlin.de

Abstract: Electric Vehicle Charging Networks (EVCNs) face critical vulnerabilities during power outages, causing transportation disruption and grid instability. Our previous work introduced EVCAR, a Multi-Agent Reinforcement Learning (MARL) framework for post-outage recovery, but it operates with limited contextual awareness. This paper presents EVCAR-KG, extending our EVCAR system with a domain-specific knowledge graph that functions as a semantic digital twin. The framework enhances our original agents with ontology-augmented perception, knowledge-guided action formulation, and semantically augmented rewards. Experimental results show EVCAR-KG achieves 20 percent faster convergence (3,013 vs 3,764 episodes) and reduces recovery time from 43.0 to 40.0 hours compared to our baseline TD3-MADDPG implementation, with only 14 percent computational overhead. The complete implementation, including the ontology definition, agent integration code, and serialized OWL file, is publicly available at: https://doi.org/10.5281/zenodo.15824931

Keywords: Electric Vehicle Charging Networks Recovery, Prolonged Service Disruption, Grid Stability, Multi-Agent Reinforcement Learning, Knowledge Graphs

DOI: 10.18154/RWTH-CONV-254920. Published in the conference proceedings of the 36. Forum Bauinformatik 2025, Aachen, Germany,

1 Introduction and Related Work

The rapid adoption of electric vehicles (EVs) presents significant challenges to the resilience of urban charging infrastructure. During power outages in densely populated districts, EV drivers often migrate en masse to neighboring areas, leading to congestion and voltage instability[1]. Events such as Berlin's 2019 blackout highlight the severity of such disruptions [2], increasing waiting times and overloading grid components.

The Electric Vehicle Charging Quality of Service Adaptive Recovery (EVCAR) system introduced a multi-agent deep reinforcement learning (MARL) approach to mitigate post-outage impacts [2]. EVCAR employs district-level agents adjusting charging rates and a central agent redistributing displaced EVs, achieving improved waiting times and grid stability compared to heuristic methods.

Building upon EVCAR, we investigate whether semantic knowledge integration can further enhance agent adaptability. Specifically, we explore ontology-driven knowledge graphs (KGs) to enrich agents' spatial, infrastructural, and operational awareness. Recent research highlights ontology-enhanced reinforcement learning (RL) for better generalization, accelerated convergence, and reliable decision-making. One direction uses ontologies to directly guide agent behavior via embedded expert knowledge. For instance, the CORL framework integrated domain ontologies into classification tasks, turning RL into structured reasoning [3]. Similar frameworks appeared in cybersecurity with ontology-defined penetration strategies combined with RL and reasoning architectures [4], [5], allowing efficient action pruning and safety constraints.

Another approach preserves RL agent autonomy while enriching observations through ontology-derived contextual knowledge. Examples include the Automatic Goal Generation Model (AGGM), dynamically generating agent goals using ontology rules during state shifts [6], and the OnCertain model enhancing RL responses during uncertainty via ontology reasoning [7]. Furthermore, ontology reasoning effectively handled sensor data gaps in smart traffic systems by inferring missing states through upstream-downstream congestion relationships [8]. Such ontologies inject semantic awareness into agents, maintaining RL exploration while improving decision-making under partial observability.

Finally, ontology-driven knowledge graphs have enhanced multi-agent coordination. Manufacturing studies used shared ontologies for real-time adaptive scheduling, dynamically aligning task and resource understanding among agents [9], [10]. This concept of shared semantic understanding is particularly relevant for decentralized multi-agent systems like EVCAR, supporting agents' coordinated efforts toward maintaining grid stability and service equity.

This paper introduces EVCAR-KG, augmenting MARL with ontology-based knowledge to assess if structured knowledge enhances training and deployment adaptability. By aligning our ontology with schemas like EVKG [11], we ensure interoperability with real-world systems.

The primary research question is: *Can integrating a domain ontology as a knowledge graph improve MARL agents' learning efficiency and decision quality in EV charging network QoS recovery through embedded contextual knowledge?*

Our contributions include:

- Designing a domain-specific ontology capturing essential EV network relationships and operational states.
- Integrating the ontology with EVCAR MARL, enabling semantic agent observations.
- Evaluating improvements in learning efficiency and QoS recovery within a controlled EVCAR environment.

2 Methodology

2.1 System Architecture Overview

EVCAR-KG extends the baseline EVCAR system through three enhancements: (1) ontology-driven knowledge graph representing EV charging network structure/state, (2) semantic query interfaces for high-level contextual information, and (3) knowledge-informed decision mechanisms including action masking and reward shaping. These extensions improve agent adaptability in variable scenarios, motivated by prior ontology-enhanced RL works [8], [5], [10]. Figure 1 shows the architecture.

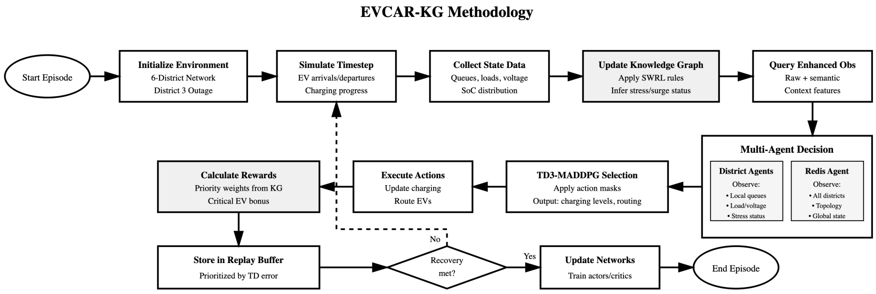

Figure 1: EVCAR-KG methodology flowchart showing the integration of ontology-driven knowledge graph with multi-agent reinforcement learning for EV charging network recovery

2.2 Ontology Architecture

The EVCAR-KG ontology provides a comprehensive semantic model of EV charging network recovery, structured in Web Ontology Language (OWL). It extends the EVKG schema [11] with domain-specific post-outage recovery concepts.

The ontology defines hierarchical classes capturing static infrastructure and dynamic operational states, distinguishing between infrastructure entities, operational events, and agent concepts (Figure 2).

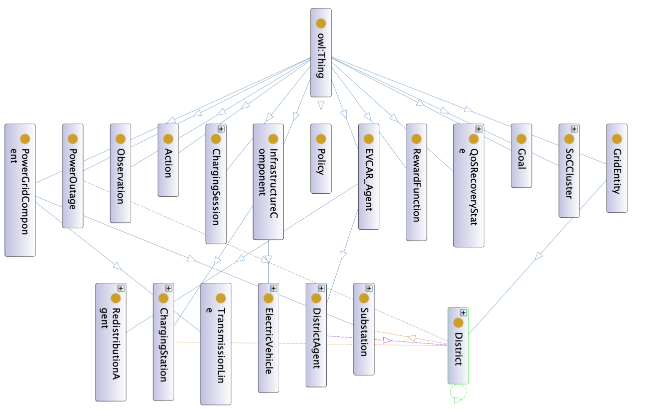

Figure 2: Ontology structure of EVCAR-KG in Protégé, showing key classes and semantic relationships relevant to EV charging network recovery.

Object properties establish network topology through spatial relationships (`locatedIn`, `hasNeighbor`), power relationships (`powers`, `suppliedBy`), and operational associations (`involvesVehicle`, `affectsDistrict`). Data properties capture both static characteristics (district types, battery capacity, station capacity) and dynamic state variables (queue lengths, voltage deviations, SoC levels, congestion states).

The ontology incorporates inference rules that derive high-level knowledge from raw data, implemented as Python functions. All rules are realised as lightweight Python callbacks; their outputs feed directly into masking and reward shaping (Table 1).

Table 1: Ontology-derived flags and their trigger conditions

Rule	Trigger condition (checked each step)	Flag/action
Grid stress	$\Delta V > 0.04$ p.u. & queue > 10 EV	`isUnderStress`=true
Surge	Neighbour outage & queue $\uparrow 50\%$ in < 1 h	`surgeStatus`=true
Priority vehicle	SoC $< 5\,\%$	`highPriority`=true (EV)
Invalid action	Route EV to stressed district	Mask / heavy penalty

The complete implementation, including the ontology definition, agent integration code, and serialized OWL file, is publicly available at: https://doi.org/10.5281/zenodo.15824931

2.3 Integration Methodology

2.3.1 State and action Augmentation

The knowledge graph fundamentally transforms how agents perceive their environment. Rather than processing raw numerical observations, agents receive semantically enriched state representations that incorporate both sensory data and inferred knowledge.

For district agents, the original observation vector from EVCAR, denoted as $o_i^d(t) \in \mathbb{R}^{15}$, undergoes augmentation with knowledge graph features to create an enhanced observation $\hat{o}_i^d(t) \in \mathbb{R}^{22}$. The augmentation process extracts seven additional features from the ontology:

$$\hat{o}_i^d(t) = [o_i^d(t), f_{type}^{res}, f_{type}^{com}, f_{type}^{ind}, f_{stress}, f_{surge}, f_{neighbors}^{out}, f_{cong}^{high}] \tag{1}$$

where f_{type}^{res}, f_{type}^{com}, and f_{type}^{ind} are binary indicators for district type (Residential, Commercial, Industrial), f_{stress} indicates whether the district is under stress according to inference rules, f_{surge} flags surge conditions, $f_{neighbors}^{out}$ counts neighboring districts with active outages, and f_{cong}^{high} indicates high congestion level.

For the redistribution agent, global state augmentation incorporates stress indicators and district types for all operational districts. The enhanced observation grows from 31 to 49 dimensions, adding three binary features per district (stress status, outage status, surge status). This comprehensive view enables system-wide coordination based on semantic understanding rather than purely numerical optimization.

The ontology also enables dynamic action space refinement through constraint propagation and safety rules. Unlike the original EVCAR where agents could select any action within their continuous or discrete space, EVCAR-KG implements ontology-based action masking.

For district agents, the action space modification occurs through a masking function $M : S \times A \to \{0, 1\}$ that filters available actions based on ontological constraints:

$$M(s,a) = \begin{cases} 0 & \text{if isUnderStress}(s) \wedge \text{increases\_load}(a) \\ 1 & \text{otherwise} \end{cases} \tag{2}$$

This masking prevents agents from accepting additional EV load when their district is already stressed, except for critical SoC vehicles which maintain priority access. The implementation modifies the actor network output by applying the mask before action selection:

$$\pi_{masked}(a|s) = \frac{\pi(a|s) \cdot M(s,a)}{\sum_{a'} \pi(a'|s) \cdot M(s,a')} \tag{3}$$

For the redistribution agent, action masking operates at the district level. Districts flagged as stressed receive reduced allocation probabilities rather than complete exclusion, maintaining system flexibility while discouraging problematic routing:

$$p_{s,k}^{masked} = \begin{cases} 0.1 \cdot p_{s,k} & \text{if isUnderStress}(k) \\ p_{s,k} & \text{otherwise} \end{cases} \tag{4}$$

This soft masking approach, where stressed districts receive 10% of their original allocation probability, prevents complete isolation while strongly discouraging additional load.

2.3.2 Reward Function Enhancement

The integration of ontological knowledge enables sophisticated reward shaping that incorporates domain expertise and operational priorities. The enhanced reward function builds upon EVCAR's multi-objective structure while adding priority-based components.

For district agents, the original reward function from EVCAR undergoes modification to incorporate SoC cluster-specific weighting:

$$\hat{r}_i^d(t) = 0.8 \left(\sum_s w_s \cdot (SC_{i,s}(t) - QL_{i,s}(t) - WT_{i,s}(t)) - L_i(t) \right) + 0.2 \cdot r_{global}(t) \tag{5}$$

where $s \in \{critical, low, medium, high\}$ represents SoC clusters, and weights follow $w_{critical} = 2.0$, $w_{low} = 1.5$, $w_{medium} = 1.0$, $w_{high} = 0.8$. This prioritization ensures that agents learn to favor critical and low SoC vehicles, improving service quality for EVs most at risk of stranding.

The global reward component incorporates penalties for violating ontology-inferred constraints:

$$r_{global}(t) = r_{global}^{base}(t) - \alpha \sum_k \mathbb{1}[\text{isUnderStress}(k)] - \beta \sum_k \mathbb{1}[\text{surgeStatus}(k)] \tag{6}$$

where α and β are penalty coefficients for stressed and surge conditions respectively. This formulation encourages system-wide stress reduction beyond local optimization.

2.3.3 Learning Process Modifications

The knowledge graph integration fundamentally alters the learning dynamics of the TD3-MADDPG algorithm. Three key modifications enhance sample efficiency and convergence properties.

Prioritized Experience Replay incorporates ontological knowledge into transition prioritization. The priority calculation extends beyond TD-error to include semantic importance:

$$p_i = |\delta_i| \cdot \omega_{semantic} \tag{7}$$

where δ_i is the TD-error and $\omega_{semantic}$ is a weighting factor based on ontological properties:

$$\omega_{semantic} = \begin{cases} 2.0 & \text{if district stressed or during outage} \\ 1.5 & \text{if critical SoC ratio} > 0.1 \\ 1.0 & \text{otherwise} \end{cases} \tag{8}$$

This prioritization ensures that agents learn more frequently from critical situations, accelerating policy improvement in high-stakes scenarios.

Results

Training comparisons show that EVCAR-KG converged at episode 3,013, approximately 20% faster than the baseline TD3-MADDPG, which required 3,764 episodes (Figure 3). This improvement is attributed to the structured domain knowledge embedded in the ontology, which supports more efficient exploration through enriched state representations. However, this gain comes with an average computational overhead of around 14%, due to increased state dimensionality and ontology-related processing. The net result is a modest reduction in total training time.

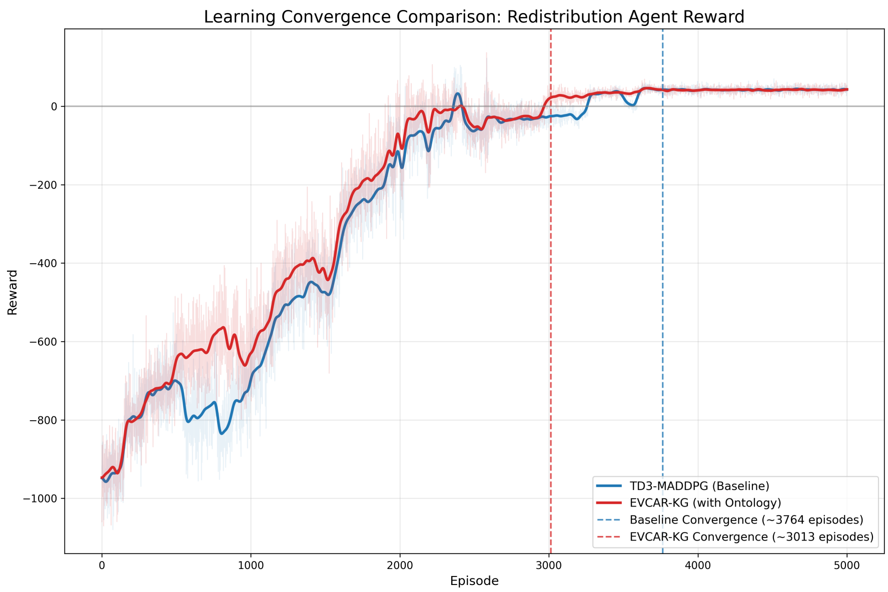

Figure 3: Comparison of agent convergence and policy evolution across episodes for base TD3-MADDPG and the integration of the ontology (EVCAR-KG).

Similarly, the district agents exhibited slightly faster convergence under EVCAR-KG, though the differences across most districts were not substantial. Districts 2 and 6, which were spatially removed from the outage, stabilized earlier, while the outage node (District 3) required more extended training, reflecting the higher complexity of local recovery under disrupted conditions.

Operationally, EVCAR-KG reduced post-outage queue lengths to acceptable levels within 40.0 hours, compared to 43.0 hours under TD3-MADDPG and significantly faster than traditional methods (60–102 hours). For vehicles in critical SoC states (<5%), both MARL approaches performed similarly, achieving safe levels within 11.9 and 12.0 hours respectively. System utilization recovered slightly faster under

EVCAR-KG, though again the margin was minimal. These results suggest that knowledge integration offers coordination advantages, particularly under spatially complex redistribution scenarios, while maintaining baseline performance across all recovery metrics.

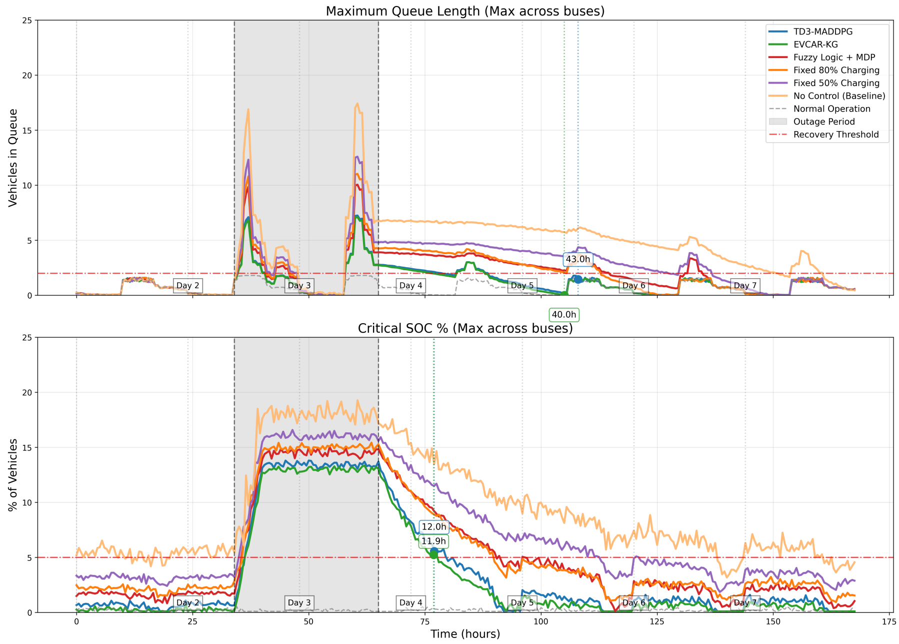

Figure 4: Comparison of user experience metrics (queue recovery and the percentage of critical SoC) across methods.

Conclusion

This paper presented EVCAR-KG, an ontology-enhanced extension of the EVCAR MARL framework for resilient recovery of electric vehicle charging networks. By integrating structured domain knowledge into the agent observation and action spaces, the system achieved slightly faster convergence and minor improvements in recovery metrics, particularly in queue length and system utilization. The performance for critical SoC management remained comparable to the baseline MARL implementation, indicating that the knowledge graph preserves coordination quality without introducing instability.

The observed gains came at the cost of approximately 14% increased computational effort. While this overhead is modest, it must be considered when deploying the system in resource-constrained environments. Nonetheless, the faster and more stable learning—particularly in complex or high-traffic districts—demonstrates the potential value of semantic knowledge integration in multi-agent recovery settings.

Future work will focus on optimizing the ontology to reduce runtime overhead, exploring the use of compressed or modular knowledge representations, and extending the system to larger urban environments with greater spatial and behavioral diversity. Additionally, real-world testing and hybrid architectures combining learning-based agents with rule-based decision modules may help strike a better balance between performance, interpretability, and scalability.

References

[1] D. Tang, X. Liu, T. Zhao, and Q. Wang, “Spatial–temporal qos assessment of the ev charging network considering power outages”, *Frontiers in Energy Research*, vol. 11, 2023. DOI: 10.3389/fenrg.2023.1112169

[2] M. Abdelfattah, N. Li, and T. Hartmann, “Evcar: Multi-agent reinforcement learning for electric vehicle charging network recovery”, in *EG-ICE 2025: International Workshop on Intelligent Computing in Engineering*, University of Strathclyde Publishing, 2025. DOI: 10.17868/strath.00093305

[3] B. Liu, L. Yao, Z. Ding, J. Xu, and J. Wu, “Combining ontology and reinforcement learning for zero-shot classification”, *Knowledge-Based Systems*, vol. 144, pp. 42–50, 2018. DOI: 10.1016/j.knosys.2017.12.022

[4] K. Qian, D. Zhang, P. Zhang, Z. Zhou, X. Chen, and S. Duan, “Ontology and reinforcement learning based intelligent agent automatic penetration test”, in *Proceedings of the 2021 IEEE International Conference on Artificial Intelligence and Computer Applications (ICAICA)*, 2021, pp. 556–561. DOI: 10.1109/ICAICA52286.2021.9497911

[5] S. Sun, Y. Lu, D. Wu, and G. Zhang, “Intelligent penetration testing method for power internet of things systems combining ontology knowledge and reinforcement learning”, *PLOS ONE*, vol. 20, no. 5, e0323357, 2025. DOI: 10.1371/journal.pone.0323357

[6] S. Ghanadbashi and F. Golpayegani, “Using ontology to guide reinforcement learning agents in unseen situations”, *Applied Intelligence*, vol. 52, no. 2, pp. 1808–1824, 2022. DOI: 10.1007/s10489-021-02449-5

[7] S. Ghanadbashi, Z. Safavifar, F. Taebi, and F. Golpayegani, “Handling uncertainty in self-adaptive systems: An ontology-based reinforcement learning model”, *Journal of Reliable Intelligent Environments*, vol. 10, no. 1, pp. 19–44, 2024. DOI: 10.1007/s40860-022-00198-x

[8] J. Guo, S. Ghanadbashi, S. Wang, and F. Golpayegani, “Urban traffic signal control at the edge: An ontology-enhanced deep reinforcement learning approach”, in *Proceedings of the 2023 IEEE International Conference on Intelligent Transportation Systems (ITSC)*, 2023, pp. 6027–6033. DOI: 10.1109/ITSC57777.2023.10422573

[9] S. Qu, J. Wang, S. Govil, and J. O. Leckie, “Optimized adaptive scheduling of a manufacturing process system with multi-skill workforce and multiple machine types: An ontology-based, multi-agent reinforcement learning approach”, *Procedia CIRP*, vol. 57, pp. 55–60, 2016. DOI: 10.1016/j.procir.2016.11.011

[10] Z. Qin and Y. Lu, “Knowledge graph-enhanced multi-agent reinforcement learning for adaptive scheduling in smart manufacturing”, *Journal of Intelligent Manufacturing*, 2024. DOI: 10.1007/s10845-024-02494-0

[11] Y. Qi, G. Mai, R. Zhu, and M. Zhang, “Evkg: An interlinked and interoperable electric vehicle knowledge graph for smart transportation system”, *Transactions in GIS*, vol. 27, no. 4, pp. 949–974, 2023. DOI: 10.1111/tgis.13064